

Robust Assignments with Vulnerable Nodes

David Adjashvili,^a Viktor Bindewald,^b Dennis Michaels^b

version from: March 20, 2017

^a Institute for Operations Research, Department of Mathematics, ETH Zurich, Rämistrasse 101, 8092 Zurich (Switzerland),
addavid@ethz.ch

^b Department of Mathematics, TU Dortmund University, Vogelpothsweg 87, 44227 Dortmund (Germany),
{viktor.bindewald, dennis.michaels}@math.tu-dortmund.de

Various real-life planning problems require making upfront decisions before all parameters of the problem have been disclosed. An important special case of such problem especially arises in scheduling and staff rostering problems, where a set of tasks needs to be assigned to an available set of resources (personnel or machines), in a way that each task is assigned to one resource, while no task is allowed to share a resource with another task. In its nominal form, the resulting computational problem reduces to the well-known assignment problem that can be modeled as matching problems on bipartite graphs.

In recent work [2], a new robust model for the assignment problem was introduced that can deal with situations in which certain resources, i.e. nodes or edges of the underlying bipartite graph, are vulnerable and may become unavailable after a solution has been chosen. In the original version from [2] the resources subject to uncertainty are the edges of the underlying bipartite graph.

In this follow-up work, we complement our previous study by considering nodes as being vulnerable, instead of edges. The goal is now to choose a minimum-cost collection of nodes such that, if any vulnerable node becomes unavailable, the remaining part of the solution still contains sufficient nodes to perform all tasks. From a practical point of view, such type of unavailability is interesting as it is typically caused e.g. by an employee's sickness, or machine failure. We present algorithms and hardness of approximation results for several variants of the problem.

Keywords. assignment problems ♦ structural robustness ♦ combinatorial optimization ♦ robust optimization ♦ approximation algorithms

1 Introduction

Input data of optimization problems is often uncertain in practice. A classical approach for dealing with uncertainty is robust optimization, an approach that seeks solutions performing well in *worst case* realizations of the uncertain parameters of the feasible set and the cost structure.

In our former work [2] a new version of Robust Assignment Problems was introduced and studied¹. The robust version studied therein is a bulk-robust problem. This robustness framework was introduced in Adjashvili et al. [1], and follows the general idea of redundancy-based robustness. Roughly speaking, bulk-robustness deals with combinatorial optimization problems in which some resources (e.g. edges or nodes in graphs) may become unavailable. This type of uncertainty is modelled by listing all subsets of resources that can simultaneously become unavailable. The union of all these subsets form the set of so-called *vulnerable* resources. Each member of the list then defines a scenario that leads to a deletion of the corresponding vulnerable resources. In this setting, a *robust solution* is a subset of resources that provides a feasible solution to the underlying nominal optimization problem in every scenario, and the task is to determine a robust solution of minimum cost.

For the Robust Assignment Problem considered in [2] the nominal optimization problem is given by the bipartite perfect matching problem while vulnerable resources are edges. We call this problem *Edge-Robust Assignment Problem* (E-RAP). The formal definition of E-RAP is the following.

Problem (The Edge-Robust Assignment Problem (E-RAP)).

- Input: Tuple (G, \mathcal{F}, c) , where $G := (U \dot{\cup} W, E)$ is a balanced, bipartite graph, i.e. $|U| = |W|$, $\mathcal{F} \subseteq 2^E$ is a family of sets of vulnerable edges, and $c \in \mathbf{R}_{\geq 0}^E$ is a non-negative cost vector.
- Output: If existent, an optimal solution to

$$\begin{aligned} \min \quad & c(X) \\ \text{s.t.} \quad & \forall F \in \mathcal{F} : X \setminus F \text{ contains a perfect matching in } G \\ & X \subseteq E. \end{aligned} \quad (\text{E-RAP})$$

The focus in [2] is on the case where $|F| = 1$ for all $F \in \mathcal{F}$. In this follow-up work, we complement our analysis for Robust Assignment Problems by now considering nodes as vulnerable resources. For this, we again consider the nominal version of the problem to be the matching problem on a bipartite graph $G = (U \dot{\cup} W, E)$, in which the goal is to match all nodes in one side of the bipartition, say U . In other words,

¹A full version of this conference contribution including all proofs is given in the preprint [3].

the nominal goal is to find a U -perfect matching. Also, we assume that all vulnerable nodes are contained in the other side of the bipartition, namely W , an assumption that is very natural the applications motivating this work. The nodes in U and W can naturally be interpreted as jobs and machines, respectively, and our assumption on the failure model implies that some machines may fail. Finally, we assume that the costs are associated with machines, i.e. with the nodes in W . Once a node $w \in W$ is included, all edges in G connecting it to nodes in U are also automatically included. Thus, a *robust solution* is a node subset $W' \subseteq W$ such that in any scenario F (corresponding to deletion of some nodes in W) the remaining nodes in $W' \setminus F$ suffice to construct a U -perfect matching, and the goal is to find a robust solution of minimum cost. We call this problem *Node-Robust Assignment Problem*, and denote it by V-RAP. Figure 1 shows an example of a V-RAP instance.




Figure 1: left: A V-RAP instance with four scenarios given by $\{1\}$, $\{2\}$, $\{4\}$, $\{6\}$ (orange nodes);
right: A node-robust solution (cyan nodes) of the V-RAP instance.

A different view on V-RAP is the following. The decision maker has to select a set of machines such that all jobs can be performed. Ahead of this decision his adversary announces a list of sets of machines from which he will sabotage one set of machines after the decision is made. The decision maker has to make sure that the act of sabotage does not jeopardize the completion of all jobs on the selected machines.

Before formally stating V-RAP, we briefly present two concrete motivating applications exhibiting its practical relevance.

- *Staff Scheduling and Rostering.* Rostering is an important task of Human Resource Management. The rostering process has four different aspects: strategic, tactical, operational and retrospective (for details we refer the reader to [9]). V-RAP addresses the tactical dimension of rostering. The objective thereby is to deliver a roster for a certain planning horizon (e.g. a month), which can be used as a starting point for operative assignment decisions on a daily (or

weekly) basis, resulting in a schedule. Naturally, a roster is subject to uncertainty because employees may get sick or not be able to work in a certain shift due to legal work load regulations or non-availability of crucial equipment. Uncertainty of this kind can be modeled in the bulk-robust framework. A solution to V-RAP enables the decision maker to choose an assignment from the solution for the upcoming day (or week), meeting the latest needs of both the company and the personnel. In the context of health care, for instance, V-RAP can be used to address continuity of service, maintaining full staff levels and minimizing costs simultaneously.

- *Subcontracting.* In a many companies numerous tasks are outsourced to subcontractors or freelancers. A typical example is the design and deployment of a new IT infrastructure. During the evaluation process a redundant set of potential subcontractors is considered. It is expected that some of the candidates will not be available after the decision process due to various reasons such as commitment elsewhere in the desired period of time or sickness of key team members. It is hence important to contact and negotiate with a redundant set of subcontractors to later be able to accommodate the latter kind of unavailabilities. V-RAP can be used to decrease the costs of the negotiation and hiring process.

The formal setting of the Node-Robust Assignment Problem as studied in this work is the following. We are given a simple bipartite graph $G = (U \dot{\cup} W, E)$ with $|U| \leq |W|$. Elements in U are associated with jobs while elements from W represent machines. In addition, each machine node w is associated with a non-negative cost $c_w \in \mathbf{R}_{\geq 0}$, yielding a cost vector $c \in \mathbf{R}_{\geq 0}^W$. A set $X \subseteq W$ of machine nodes is called an *assignment* of G if the subgraph $G[U \dot{\cup} X]$ induced by the selected machine nodes in X and all job nodes from U contains a U -perfect matching, i.e. a set of non-adjacent edges incident to all job nodes in U . In the nominal version the objective is to find an assignment of G with minimum cost. In the node-robust version, the set W of machine nodes is subject to uncertainty. The uncertainty is described by a list $\mathcal{F} \subseteq 2^W$ of subsets of machine nodes. Each element F of \mathcal{F} defines a failure scenario, the occurrence of which involves the removal of all machine nodes in F from G . We are now interested in determining a subset $X \subseteq W$ of machine nodes in such a way that, for every failure scenario $F \in \mathcal{F}$, the induced subgraph $G[U \dot{\cup} X] - F$ contains a U -perfect matching. We call such a subset a *node-robust solution*. The aim of V-RAP is to find, for a bipartite graph G and for a set \mathcal{F} of failure scenarios, a node-robust solution that is minimal with respect to the cost vector c .

We will restrict ourselves in this manuscript to the case when each failure scenario is defined by one single machine node, i.e. we assume that \mathcal{F} is a family of singletons, and hence can simply be represented a subset of W , namely $\mathcal{F} \subseteq W$. As shown in

our previous work for E-RAP, we will prove in this work that the node-robust version exhibits interesting structure under this assumption. In particular, we will show that this case is hard to approximate within a sub-logarithmic factor in general, and within a constant greater than one in the unweighted case. We also provide approximation algorithms with matching approximation guarantees (up to small additive constants).

The Node-Robust Assignment Problem under consideration is formally stated next.

Problem (The Node-Robust Assignment Problem (V-RAP)).

- Input: Tuple (G, \mathcal{F}, c) , where $G := (U \dot{\cup} W, E)$ is a bipartite graph, $\mathcal{F} \subseteq W$ is a set of vulnerable nodes, and $c \in \mathbf{R}_{\geq 0}^W$ is a non-negative cost vector.
- Output: If existent, an optimal solution to

$$\begin{aligned} \min \quad & c(X) \\ \text{s.t.} \quad & \forall f \in \mathcal{F} : G[U \dot{\cup} X] - \{f\} \text{ contains a } U\text{-perfect matching,} \\ & X \subseteq W. \end{aligned} \quad (\text{V-RAP})$$

There are two important cases of V-RAP that are worth mentioning. The first case concerns the structure of the uncertainty set. A V-RAP instance $((U \dot{\cup} W, E), \mathcal{F}, c)$ is called *uniform* if each machine node is vulnerable, i.e. $\mathcal{F} = W$. The second case is that of unit costs, i.e. the case $c_w = 1$, for each $w \in W$. In this case, one is interested in finding a node-robust solution with minimum cardinality. We will denote this case by card-V-RAP.

We remark that it is easy to verify whether a given instance (G, \mathcal{F}, c) of V-RAP is feasible. This can be achieved by applying any efficient bipartite maximum matching algorithm $|\mathcal{F}| \leq |W|$ times. Therefore, we will assume that every V-RAP instance considered in this work is feasible.

Analogously to [2] for the Edge-Robust Assignment Problem, we focus in this work on deriving hardness results and approximation algorithms for V-RAP and card-V-RAP. Our results are summarized in Table 1. This table also compares our results for V-RAP with our results for E-RAP.

Problem ($\mathcal{F} \subseteq W, E$)	hardness of approximation	algorithm's guarantee
V-RAP	$d \log n, d < 1$ [Thm. 2]	$\log n + 2$ [Thm. 5]
card-V-RAP	no PTAS [Thm. 6]	1.75 [Thm. 7] (2 on non-uniform instances)
E-RAP	$d \log n, d < 1$ [2, Thm. 3]	$O(\log n)$ (randomized) [2, Thm. 4]
card-E-RAP [¶]	no PTAS [2, Thm. 5]	1.5 [2, Thm. 6] (3 on non-uniform instances)

Table 1: Summary of results for V-RAP as will be presented in this work and for E-RAP as stated in [2] (for proofs, see the full preprint version [3]) where n represents the number of nodes in the underlying bipartite graph.

[¶] card-E-RAP denotes the unweighted Edge-Robust Assignment Problem.

The paper is organized as follows. In Section 2 we briefly review existing results related to V-RAP (for related work concerning E-RAP see [2]). Section 3 deals with the general, weighted version of V-RAP, for which we show hardness of approximation and present an approximation algorithm with an approximation. The bounds depend logarithmically on the number of nodes. In Section 4 we focus on the unweighted version card-V-RAP. We first prove that there cannot exist a PTAS for card-V-RAP, unless $P = NP$. We then analyze the approximation algorithm introduced for general V-RAP for the unweighted case. We show that this algorithm becomes a constant factor approximation algorithms for card-V-RAP.

One interesting fact about the Edge-Robust Assignment Problem is that this problem is even NP-hard in its simplest variant, i.e. in case of two vulnerable edges and unit weights (see [2], and [3, Section 5] for a proof). In Section 4.3 we show that, in contrast to the situation E-RAP, the case of two vulnerable machines nodes is solvable in polynomial time.

Notation. Throughout this work we use the following notation. Let G be a bipartite graph. By $V(G)$ and $E(G)$ we denote the node set and the edge set of G . For a subset $E' \subseteq E(G)$, $V(E')$ is used to represent the set of all nodes incident by E' . For a subset $V' \subseteq V(G)$, $E(V')$ denotes the set of edges of G with both endpoints in V' . The subgraph $(V', E(V'))$ induced by a node subset V' is abbreviated by $G[V']$. For a subset $V' \subseteq V(G)$ of nodes, we use the notation $G - V'$ to denote the graph resulting from G when all nodes from V' and all edges incident to some node of V' are removed.

2 Related work

Laroche et al. [14] considered a problem related to V-RAP. For a given bipartite graph $G := (U \dot{\cup} W, E)$ they consider the following interdiction problem: Does the removal of k arbitrary nodes from W results in a graph without a U -perfect matching? The authors were especially interested in computing the smallest number k for which the answer to the latter question is yes. The motivation to study the problem comes from the nurse rostering problem arising in health care. In that context one is interested in determining the largest number of nurses that can be absent such that all patients can still be treated adequately. This largest number can be seen as a measure for the resilience of a health care provider with respect to staff unavailability. In the setting of V-RAP, the question above can be interpreted as follows. Given a bipartite graph $G = (U \dot{\cup} W, E)$, what is the largest k such that the machine node set W is a feasible solution for V-RAP when the scenarios are defined by $\mathcal{F}_k = \{F \subseteq W : |F| = k\}$. To answer the question, the authors exploit in [14] the so-called k -extended Hall's condition:

$$\forall \emptyset \neq T \subseteq U : |T| + k \leq |N_G(T)| \quad (\text{k-Hall})$$

which guarantees the existence of a node-robust assignment w.r.t. the uncertainty set \mathcal{F}_k . Because G is bipartite, this question can be answered efficiently by solving an ILP over an integral polytope $|U|$ times.

Remark. The results from Laroche et al. [14, Cor. 2] imply that feasibility testing for V-RAP can be performed in polynomial time when the list \mathcal{F} of scenarios is given implicitly. This is a major difference to E-RAP where deciding upon feasibility of an instance with implicitly given scenarios is an NP-complete problem (cf. [3]), which is an immediate consequence of the NP-completeness of the so-called Matching Preclusion Number Problem (Lacroix et al. [13, Thm. 6], Dourado et al. [8, Thm. 2])

Zenklusen [17] considered the Matching Interdiction problem that asks for a given node-weighted graph and a budget to find a subset of nodes that respects the budget constraint and that minimizes the size of a maximum matching when the selected nodes are removed from the graph.

Arulselvan et al. [6] analyzed the following variant of assignment problems. The input consists of an edge-weighted bipartite graph $G = (U \dot{\cup} W, E)$ and lower and upper quotas $l, u \in \mathbf{Z}_{\geq 0}^W$. The goal is to find a maximum-weight edge set M such that each node in U is incident to at most one edge in M . Furthermore, the nodes in W have either to respect the bounds given by the quota functions or not to be used at all. Their main results yield a classification of several variants of the problem in terms of their complexity.

Katriel et al. [12] studied a two-stage stochastic optimization problem on a bipartite graph $G := (U \dot{\cup} W, E)$ with a cost function on the edge set, that resembles V-RAP.

The overall optimization task is to compute an edge subset that contains a maximum matching. In the first stage there is no uncertainty and one can already select some edges at nominal costs. In the second stage, uncertainty comes into play in two variants: either the costs on the edges are uncertain, or some of the nodes from W are deactivated. For both variants the goal in the second stage is to buy additional edges such that the edges bought in the two stages contain a maximum matching at minimum expected costs. The main results include the derivation of lower bounds on approximation guarantees as well as approximation algorithms. Furthermore, the authors provide a randomized approximation algorithm for the robust variant of the stochastic optimization problem under consideration.

The connection between the matching number of a graph and node removal is also studied from a graph theoretical point of view. For instance, Aldred et al. [4] provided conditions under which special graph classes (grid graphs and k -fold product graphs) remain perfectly matchable after node deletions. Favaron [10] investigated and characterized the class of so-called k -factor-critical graphs, i.e. graphs on n nodes such that every subgraph on $n - k$ nodes is perfectly matchable. Note that such graphs can not be bipartite, hence those insights can not be applied here.

3 Hardness and approximability of V-RAP

In this section we show that V-RAP is hard to approximate within a factor of $d \log n$ for any $d < 1$ and provide an approximation algorithm with a matching approximation ratio (up to an additive constant)

The hardness of approximation is derived in Section 3.1. As in the case for the Edge-Robust Assignment Problem [2, 3], the key ingredient is a reduction from the set cover problem.

The approximation algorithm for V-RAP is presented and discussed in Section 3.2. The core idea of the algorithm is the following. The algorithm first determines a certain matching and selects all machine nodes covered by the matching. In order to extend this set of machines to a node-robust solution, a set cover instance is constructed on all job nodes not saturated by the machine nodes selected through the matching. The set cover instance is then solved with the greedy algorithm [7].

3.1 Hardness for V-RAP

To prove our hardness result for V-RAP, we present a reduction from the NP-hard set cover problem.

Problem (Set Cover Problem (SC)).

- Input: Tuple $([k], \mathcal{S})$ with a finite ground set $[k] = \{1, \dots, k\}$ and a collection $\mathcal{S} := \{S_1, \dots, S_l\}$ of subsets of $[k]$, for some $k, l \in \mathbf{Z}_{\geq 1}$.
- Output: Collection $\mathcal{C} \subseteq \{S_1, \dots, S_l\}$ with $\bigcup_{S \in \mathcal{C}} S = [k]$ minimizing $|\mathcal{C}|$.

For a given instance $([k], \mathcal{S})$ of SC, existence of any cover for ground set $[k]$ can be efficiently verified, simply by checking if $\bigcup_{S \in \mathcal{S}} S = [k]$ holds. Thus, we will assume from now on that any SC instance considered is feasible.

Now let $\mathcal{I} := ([k], \mathcal{S})$ be any feasible instance of SC. We associate with \mathcal{I} the following instance $\mathcal{I}' := (G, \mathcal{F}, c)$ of V-RAP. The graph G is obtained by applying the following steps.

- (T1) For each $s \in [k]$, node u_s is introduced and added to $U_{[k]}$. For each $S_j \in \mathcal{S}$, node w_{S_j} is introduced and added to $W_{\mathcal{S}}$. Furthermore, the edge $\{u_s, w_{S_j}\}$ is introduced and added to E_{SC} whenever $s \in S_j$.
- (T2) For each $s \in [k]$, a copy w_s of u_s is introduced and added to $W_{[k]}$, and the edge $\{u_s, w_s\}$ is added to $E_{[k]}$.

Note that edges from E_{SC} encode whether an element $s \in [k]$ is contained in a subset $S \in \mathcal{S}$, or not. Nodes from $W_{[k]}$ are used to ensure the feasibility of the V-RAP instance, while the nodes from $W_{\mathcal{S}}$ indicate which elements from \mathcal{S} are chosen to cover the ground set $[k]$.

Applying steps (T1) and (T2) yields the graph $G := (U_{[k]} \dot{\cup} (W_{\mathcal{S}} \cup W_{[k]}), E_{\text{SC}} \cup E_{[k]})$. An example of such a graph is illustrated in Figure 2.




Figure 2: Graph G corresponding to a Set Cover instance with ground set $[k]$ and covering sets S_1, \dots, S_l .

To complete the construction of \mathcal{I}' we set $\mathcal{F} = W_{[k]}$ and

$$c \in \mathbf{R}_{\geq 0}^{W_{\mathcal{S}} \cup W_{[k]}} \quad \text{with} \quad w \mapsto c_w := \begin{cases} 1, & \text{if } w \in W_{\mathcal{S}}, \\ 0, & \text{if } w \in W_{[k]}. \end{cases} \quad (1)$$

The latter transformation can clearly be carried out in polynomial time.

The next lemma highlights the relation between the SC and V-RAP instances.

Lemma 1. *Let $\mathcal{I} := ([k], \mathcal{S})$ be a (feasible) instance of SC, and let $\mathcal{I}' := (G, \mathcal{F}, c)$ be the corresponding V-RAP instance with $G := (U_{[k]} \dot{\cup} (W_{\mathcal{S}} \cup W_{[k]}), E_{[k]} \cup E_{\text{SC}})$ obtained by applying steps (T1) – (T2), uncertainty set $\mathcal{F} = W_{[k]}$ and cost vector $c \in \mathbf{R}_{\geq 0}^{W_{\mathcal{S}} \cup W_{[k]}}$ as specified in Equation (1). Then, for $X \subseteq W_{\mathcal{S}} \cup W_{[k]}$ with $W_{[k]} \subseteq X$, it holds that X is feasible to \mathcal{I}' if and only if $\mathcal{C}_X := \{S_j \in \mathcal{S} \mid w_{S_j} \in X\}$ is feasible to \mathcal{I} . Furthermore, such sets X and \mathcal{C}_X have identical costs.*

Proof. Let $\mathcal{I} := ([k], \mathcal{S})$ be a given SC instance, and let $\mathcal{I}' := (G, W_{[k]}, c)$ be the corresponding V-RAP instance. “only if” part. Let X be any feasible solution to \mathcal{I}' with $W_{[k]} \subseteq X$, and let $s \in [k]$. To show that s is contained in some set of \mathcal{C}_X , consider the node $w_s \in \mathcal{F}$. As X is feasible to \mathcal{I}' , there must exist a $U_{[k]}$ -perfect matching M in $G[U_{[k]} \dot{\cup} X] - \{w_s\}$. Since $w_s \notin V(M)$, node u_s must be matched with some node from $\{w_{S_1}, \dots, w_{S_l}\}$ by the corresponding edge from E_{SC} , i.e. $w_{S_j} \in X$, for some $S_j \in \mathcal{S}$ with $s \in S_j$. This implies that $S_j \in \mathcal{C}_X$. It follows that s is covered, and that \mathcal{C}_X is a feasible cover for \mathcal{I} .

“if” part. Let $\mathcal{C} \subseteq \mathcal{S}$ be a feasible cover for \mathcal{I} . Then, define

$$X := W_{[k]} \cup \{w_{S_j} \mid S_j \in \mathcal{C}\}, \text{ implying that } \mathcal{C} = \mathcal{C}_X \text{ holds.}$$

Recall that X is feasible to \mathcal{I}' if and only if $X \setminus \{w_s\}$ contains an assignment of G , for all $w_s \in \mathcal{F} = W_{[k]}$. Consider an arbitrary $w_s \in \mathcal{F}$. We have to show that $G[U_{[k]} \dot{\cup} X] - \{w_s\}$ contains a $U_{[k]}$ -perfect matching. Let $S_j \in \mathcal{C}$ be any set covering s . Such a set exists, since \mathcal{C} is a cover. The desired perfect matching can now be defined as

$$M := \{u_s, w_{S_j}\} \cup E_{[k]} \setminus \{\{w_s, u_s\}\}.$$

Finally, the costs of the two solutions is clearly identical by the definition of the reduction. \blacksquare

Theorem 2. *Unless $\text{NP} \subseteq \text{DTIME}(n^{\log \log n})$, V-RAP admits no polynomial $d \log n$ -approximation algorithm for any $d < 1$, where n represents the number of nodes in the underlying graph.*

Proof. Feige [11] showed that, for any $d < 1$, SC admits no polynomial time $d \log n$ -approximation algorithm (n being the size of the ground set) unless we have that $\text{NP} \subseteq \text{DTIME}(n^{\log \log n})$. Combining this result with Lemma 1 yields the proof. ■

We conclude by showing that the result in Theorem 2 also holds the more restricted uniform case, in which every machine is vulnerable.

Proposition 3. *Theorem 2 even holds for the uniform case, i.e. with $\mathcal{F} = W_S \cup W_{[k]}$.*

Proof. The proof follows from the same reduction with the only difference that the set of vulnerable machine nodes contains all machine nodes (resulting in a uniform instance). The fact that all machines in $W_{[k]}$ have cost zero implies that we can assume them to be part of any solution. The inclusion of these machines allows to match all jobs in the case that any machine in W_S fails. The rest of the proof remains the same. ■

3.2 A $(\log |U| + 2)$ -approximation for V-RAP

In this section we present a $(\log |U| + 2)$ -approximation algorithm for V-RAP stated formally as Algorithm 1. For a node $u \in U \cup W$ we denote by $\delta(u) \subseteq E$ the set of edges incident to u .

The algorithm starts by computing a U -perfect matching M (step 2) minimizing the cost of job nodes W^M that are incident to the matching. This is easily achieved by defining an appropriate cost function on the edge set and computing a minimum cost perfect matching. The set W^M is chosen to be part of the solution (step 3) and then extended to a feasible V-RAP solution as follows. The algorithm identifies all job nodes that are matched by M with a vulnerable machine node in W^M , yielding the set $U_{\mathcal{F}}$ (step 4) called *critical nodes*. We think of the matching M as the basis of any U -perfect matching for each of the failure scenarios, and the additional machine nodes we add are designed to replace some vulnerable machine in W^M , in case it fails. The set $U_{\mathcal{F}}$ is hence the set of jobs that may become unmatched by removing from M edges that are incident to vulnerable nodes.

When a node in $u \in U_{\mathcal{F}}$ becomes unmatched in M due to failure of its corresponding machine $w \in W^M$, it is possible to obtain a new U -perfect matching by finding an M -alternating path starting at u , ending with some node $w' \in W \setminus W^M$ and not using the edge $\{u, w\}$. Such a path starts and ends with an edge not in M , and hence it can be used to increase the matching size by one, resulting in a new U -perfect matching. To allow this new matching we only need to include the end-node w' of this path in the solution. We can hence think of w' as a node in $W \setminus W^M$ covering the scenario corresponding to u .

This covering interpretation is then used in the algorithm to define a set covering problem. For each machine node $w \in W \setminus W^M$, the algorithm determines a subset $R_w \subseteq U_{\mathcal{F}}$ of critical job nodes u (steps 5 – 7) with the property that there is an M -alternating u - w path in G . Finally, a weighted set cover instance \mathcal{I}^{SC} with ground set $U_{\mathcal{F}}$ and with the collection of subsets R_w , $w \in W \setminus W^M$, is constructed (step 8) and approximately solved by the greedy algorithm (step 9).

Algorithm 1 : A $(\log |U| + 2)$ -approximation for V-RAP

Require: A feasible V-RAP-instance \mathcal{I} : $G = (U \dot{\cup} W, E)$, $c \in \mathbf{R}_{>0}^W$, $\mathcal{F} \subseteq W$.

Ensure: A feasible solution X .

- 1: Define an auxiliary cost function $d \in \mathbf{R}^E$: for each node $w \in W$ and for each $e \in \delta(w)$ set $d_e := c_w$.
 - 2: $M \leftarrow$ minimum-cost U -perfect matching w.r.t. cost function d
 - 3: $W^M \leftarrow V(M) \cap W$
 - 4: $U_{\mathcal{F}} \leftarrow \{u \in U \mid u \text{ is matched to a vulnerable machine node in } M\}$
 - 5: **for** each machine node $w \in W \setminus W^M$ **do**
 - 6: $R_w \leftarrow \{u \in U_{\mathcal{F}} \mid G[U \dot{\cup} (W^M \cup \{w\})] \text{ contains an } M\text{-alternating } u\text{-}w\text{-path}\}$
 - 7: **end for**
 - 8: Construct a set cover instance $\mathcal{I}^{\text{SC}} := (U_{\mathcal{F}}, \{R_w \mid w \in W \setminus W^M\})$ with weight function $g : \{R_w \mid w \in W \setminus W^M\} \rightarrow \mathbf{R}_{\geq 0}$, $R_w \mapsto g(R_w) := c_w$
 - 9: Apply the greedy algorithm to the set cover instance \mathcal{I}^{SC} to obtain an approximate solution W^{SC} .
 - 10: **return** $X = W^M \dot{\cup} W^{\text{SC}}$
-

As we informally sketched above, the approximate solution of \mathcal{I}^{SC} defines a set W^{SC} of machine nodes whose addition to W^M results in a feasible solution X for the V-RAP instance. This property is established by the next lemma.

Lemma 4. *Let $\mathcal{I} = (G, \mathcal{F}, c)$ be a feasible V-RAP instance with $G := (U \dot{\cup} W, E)$, and let M be any U -perfect matching in G . We denote by W^M the machine nodes covered by M . Furthermore, let $U_{\mathcal{F}}$ be the set of all critical job nodes as determined in step 4 of Algorithm 1, and let $X \subseteq W$ be any subset of machine nodes containing W^M . Define the set system $\mathcal{R}_X^{\text{SC}} := \{R_w \mid w \in X \setminus W^M\}$ where R_w represents the set of critical job nodes calculated in step 6 of Algorithm 1.*

Then, X is feasible for \mathcal{I} if and only if $\mathcal{R}_X^{\text{SC}}$ forms a cover for $U_{\mathcal{F}}$ (i.e. $\mathcal{R}_X^{\text{SC}}$ is a feasible solution for the set cover instance \mathcal{I}^{SC} constructed in step 8 of Algorithm 1).

Proof. Let $X \subseteq W$ with $W^M \subseteq X$ be a feasible solution to \mathcal{I} , and consider an arbitrary node $u \in U_{\mathcal{F}}$. We have to show that there exists a set R_w containing u , for some $w \in X \setminus W^M$.

Now, let f be the vulnerable machine matched to u by M . We consider the matching $\hat{M} := M \setminus \{\{u, f\}\}$ of size $|U| - 1$. As X is feasible, the subgraph $H := G[U \dot{\cup} X] - \{f\}$

contains a U -perfect matching that has size $|U|$. Since \hat{M} is also a matching in H , it follows from a result by Berge that it can be augmented to a U -perfect matching using an \hat{M} -augmenting path \hat{P} . The path \hat{P} starts in u and ends in some node $w \in X \setminus W^M$. Since \hat{P} is \hat{M} -augmenting, w is the only node in \hat{P} that is not incident to \hat{M} , i.e. \hat{P} is a path in $G[U \dot{\cup} (W^M \cup \{w\})]$. By construction of \hat{M} , \hat{P} must be an odd M -alternating path and hence $u \in R_w$. As $u \in U_{\mathcal{F}}$ was chosen arbitrarily, we can conclude that $\mathcal{R}_X^{\text{SC}}$ forms a cover of $U_{\mathcal{F}}$ and is hence feasible to \mathcal{I}^{SC} .

For the reverse direction, let $\mathcal{R}^{\text{SC}} \subseteq \{R_w \mid w \in W \setminus W^M\}$ be any feasible solution to the set cover instance \mathcal{I}^{SC} as constructed in step 8 of Algorithm 1. We have to show that $X = W^M \dot{\cup} \{w \in W \setminus W^M \mid R_w \in \mathcal{R}^{\text{SC}}\}$ is feasible for \mathcal{I} . For this, we prove that for each vulnerable machine node $f \in \mathcal{F}$, there exists a U -perfect matching not using f in the graph $G[U \dot{\cup} X]$. Note that M is contained in $G[U \dot{\cup} X]$, i.e. M provides the desired U -perfect matching for all $f \in \mathcal{F}$ not incident to M .

Now, consider an arbitrary $f \in \mathcal{F}$ incident to M . We denote by u the job node matched to f in M . By definition, u is a critical job node, i.e. $u \in U_{\mathcal{F}}$. From the fact that \mathcal{R}^{SC} forms a cover of $U_{\mathcal{F}}$, it follows that $u \in R_w$, for some $w \in X \setminus W^M$. By definition of R_w , there exists an M -alternating path P from u to w in $G[U \dot{\cup} (W^M \cup \{w\})]$. As w is not covered by M , the path P ends with an edge (incident to w) that does not belong to M . Note that P has an odd number of edges. Thus, the first edge of P incident with u is also not contained in M , i.e. $\{u, f\} \notin P$. Thus, $M \triangle (P + \{u, f\})$ is a U -perfect matching in $G[U \dot{\cup} X]$ not using f . ■

We are now ready to prove the main result of this section.

Theorem 5. *Algorithm 1 is a polynomial $(\log |U| + 2)$ -approximation algorithm for V-RAP.*

Proof. The algorithm can clearly be implemented in polynomial time. For a reference on algorithms for minimum cost perfect matching computations and augmenting path computations we refer the reader to [15].

Let \mathcal{R}^{SC} be the approximate solution of the set cover instance computed by the greedy algorithm. Then, Algorithm 1 returns $X = W^M \dot{\cup} W^{\text{SC}}$ where $W^{\text{SC}} := \{w \in W \setminus W^M \mid R_w \in \mathcal{R}^{\text{SC}}\}$. As \mathcal{R}^{SC} forms a cover of $U_{\mathcal{F}}$, it follows from Lemma 4 that X is feasible for \mathcal{I} .

It remains to prove that the computed solution X satisfies the desired quality. For this, let OPT and OPT^{SC} be optimal solutions for the given V-RAP instance \mathcal{I} and for the constructed set cover instance \mathcal{I}^{SC} with associated cost $c(\text{OPT})$ and $g(\text{OPT}^{\text{SC}}) = \sum_{R_w \in \text{OPT}^{\text{SC}}} g(R_w) = \sum_{R_w \in \text{OPT}^{\text{SC}}} c_w$. The cost of the returned solution X is

$$c(X) = c(W^M) + c(W^{\text{SC}}). \quad (2)$$

As M is chosen in step 2 to be a U -perfect matching minimizing the cost of the incident machine nodes and since any feasible V-RAP solution contains some U -perfect matching we obtain

$$c(W^M) \leq c(\text{OPT}). \quad (3)$$

Since we solve the set cover instance with the greedy algorithm we have

$$c(\mathcal{R}^{\text{SC}}) \leq (\log |U_{\mathcal{F}}| + 1) \cdot c(\text{OPT}^{\text{SC}}), \quad (4)$$

due to [7]. As OPT is feasible to the V-RAP instance \mathcal{I} , Lemma 4 implies that the associated collection $\mathcal{R}_{\text{OPT}}^{\text{SC}} = \{R_w \mid w \in \text{OPT} \setminus W^M\}$ is feasible to the set cover instance \mathcal{I}^{SC} . Thus, we further have that

$$g(\text{OPT}^{\text{SC}}) \leq g(\mathcal{R}_{\text{OPT}}^{\text{SC}}) = \sum_{R_w \in \mathcal{R}_{\text{OPT}}^{\text{SC}}} c_w = c(\text{OPT} \setminus W^M) \leq c(\text{OPT}). \quad (5)$$

Combining the results from Equations (2)–(5) and using $U_{\mathcal{F}} \subseteq U$, we obtain the desired bound on the approximation guarantee

$$c(X) \leq (\log |U| + 2) \cdot c(\text{OPT}).$$

■

We note that the approximation guarantee we obtain is actually $\log |U_{\mathcal{F}}| + 2$, which is at most

$$\log \min\{|U|, |\mathcal{F}|\} + 2,$$

since $|U_{\mathcal{F}}| \leq \min\{|U|, |\mathcal{F}|\}$ by definition of the matching M and $U_{\mathcal{F}}$. We hence obtain a better guarantee if the number of faulty machine nodes is significantly smaller than $|U|$.

4 Hardness and approximability of card-V-RAP

This section deals with the unweighted version card-V-RAP. Similarly to the unweighted edge-robust variant card-E-RAP, we first prove in Section 4.1 that card-V-RAP does not admit a PTAS, provided that $\text{P} \neq \text{NP}$. This is achieved by reducing the vertex cover problem in sub-cubic graphs to card-V-RAP and invoking a result of Alimonti and Kann [5] for the former problem. In Section 4.2, we refine the analysis of our approximation algorithm presented for general V-RAP. In particular, we show that this algorithm is a constant factor approximation algorithm for card-V-RAP instances. In Section 4.3 we consider the case when the number of vulnerable machine nodes is two. We prove that card-V-RAP is solvable in polynomial time in this case.

4.1 Hardness of card-V-RAP

In this section we prove the following result for card-V-RAP.

Theorem 6. *For some constant $\delta > 1$, there is no polynomial δ -approximation for uniform card-V-RAP, unless $P = NP$.*

To prove Theorem 6, we resort to the reduction from Set Cover discussed in Section 3. Since we now wish to deal with uniform weights on the machine nodes, the former reduction needs to be adjusted. In particular, we cannot assume that $W_{[k]}$ is contained in some optimal solution since including all of $W_{[k]}$ might always be sub-optimal. To be more precise, let $\mathcal{I} := ([k], \mathcal{S})$ be a feasible instance of SC, and let $G := (U_{[k]} \dot{\cup} (W_{\mathcal{S}} \cup W_{[k]}), E_{\text{SC}} \cup E_{[k]})$ be the resulting graph obtained by performing steps (T1) and (T2) from Section 3. We apply the following additional step to G .

(T3) For each $s \in [k]$, two further copies \bar{u}_s and \bar{w}_s are introduced and added to $\bar{U}_{[k]}$ and to $\bar{W}_{[k]}$, respectively. Then, for each $s \in [k]$, the edges $\{w_s, \bar{u}_s\}$ and $\{\bar{u}_s, \bar{w}_s\}$ are introduced and added to $E_{[k]}$.

Let $\bar{G} := (\bar{U} \dot{\cup} \bar{W}, \bar{E})$ be the resulting graph with $\bar{U} := U_{[k]} \cup \bar{U}_{[k]}$, $\bar{W} := W_{\mathcal{S}} \cup W_{[k]} \cup \bar{W}_{[k]}$ and $\bar{E} := E_{\text{SC}} \cup E_{[k]}$. An example of such a graph is illustrated in Figure 3. Note that the size of \bar{G} is still polynomial in k and l , the input length of the corresponding set cover instance \mathcal{I} . We can now prove Theorem 6.




Figure 3: The graph \bar{G} corresponding to a Set Cover instance with ground set $[k]$ and covering sets S_1, \dots, S_l .

Proof. It is well known that the *Vertex Cover Problem in sub-cubic Graphs* (VC3) on an input graph $H = (V_H, E_H)$ with $|E_H| = k$ can be equivalently restated as

an instance of the Set Cover Problem, where the ground set is $E_H \cong [k]$ and each $S_j \in \mathcal{S}$ corresponds to a cut set $\delta(v)$, for some $v \in V_H$, i.e. $\mathcal{S} = \{\delta(v) \mid v \in V_H\}$. As H is sub-cubic, we have that $|S_j| \leq 3$, for all $S_j \in \mathcal{S}$. Moreover, $|\{S_j \in \mathcal{S} \mid s \in S_j\}| = 2$ holds for each $s \in [k]$, namely $s \in S_{j_1}$ and $s \in S_{j_2}$ where s represents edge $e = \{v_1, v_2\} \in E_H$ and S_{j_1} and S_{j_2} correspond to $\delta(v_1)$ and $\delta(v_2)$.

Now, let a VC3 instance be presented as a SC instance $([k], \mathcal{S})$. Furthermore, let \bar{G} be the graph obtained from $([k], \mathcal{S})$ by applying the transformation steps (T1)–(T3), and let $\bar{\mathcal{I}}$ be the uniform card-V-RAP instance induced by \bar{G} .

Lemma 1 has to be adjusted in the following manner. First note that any feasible solution $X \subseteq \bar{W}$ to V-RAP with $\mathcal{F} = \bar{W}$ has to contain both sets $W_{[k]}$ and $\bar{W}_{[k]}$, because the jobs corresponding to node set $\bar{U}_{[k]}$ are only adjacent to those nodes. The vulnerability of nodes in $W_{[k]}$ implies again that X corresponds to the feasible cover $\mathcal{C}_X := \{S_j \mid w_{S_j} \in X\}$. Conversely, any feasible cover \mathcal{C} for $([k], \mathcal{S})$ gives rise to a feasible solution $X_{\mathcal{C}} := \{w_{S_j} \mid S_j \in \mathcal{C}\} \cup W_{[k]} \cup \bar{W}_{[k]}$ for the corresponding uniform card-V-RAP instance with cost $|X_{\mathcal{C}}| = 2k + |\mathcal{C}|$.

Next, we observe that \bar{G} contains $5k$ edges (as $([k], \mathcal{S})$ encodes a VC3 instance) and that any feasible solution of VC3 has size of at least $k/3$. Now assume on contrary that for any constant $\delta > 1$, there is a polynomial δ -approximation A_{δ} for uniform card-V-RAP. Due to the relation $|X_{\mathcal{C}}| = 2k + |\mathcal{C}|$, A_{δ} can be used to approximate VC3 within a certain constant factor $\alpha(\delta) > 1$ where $\alpha(\delta) \rightarrow 1$, for $\delta \rightarrow 1$. However, Alimonti and Kann proved in [5] that there exists a constant $\alpha > 1$ such that VC3 does not admit a polynomial α -approximation algorithm unless $P=NP$. This completes the proof. \blacksquare

4.2 A $O(1)$ -approximation algorithm for uniform card-V-RAP

In this section we prove the following result.

Theorem 7. *Uniform card-V-RAP admits a polynomial 1.75-approximation algorithm.*

Proof. We show that Algorithm 1, when applied to feasible uniform card-V-RAP instances, already guarantees the desired approximation ratio. For clarity we rewrite Algorithm 1 in terms of uniform card-V-RAP to obtain Algorithm 2. Since we aim to obtain a constant factor approximation ratio we can no longer directly rely on the approximation guarantee of the greedy algorithm of an arbitrary set cover instance, and instead perform an analysis of the entire algorithm simultaneously. We thus write the steps of the greedy algorithm explicitly in Algorithm 2.

Algorithm 2 : A 1.75-approximation algorithm for uniform card-V-RAP.

Require: A feasible card-V-RAP-instance: $G = (U \dot{\cup} W, E)$, $\mathcal{F} = W$.

Ensure: A feasible solution X .

```

1:  $M \leftarrow$  any  $U$ -perfect matching in  $G$ 
2:  $W^M \leftarrow V(M) \cap W$ 
3: for each machine node  $w \in W \setminus X$  do
4:    $R_w \leftarrow \{u \in U \mid \exists \text{ an } M\text{-alternating } u\text{-}w\text{-path in } G[U \dot{\cup} (W^M \cup \{w\})]\}$ 
5: end for
6:  $X \leftarrow W^M$ 
7:  $U_{\mathcal{F}} \leftarrow U$ 
8: while  $|U_{\mathcal{F}}| > 0$  do
9:    $\bar{w} \leftarrow \arg \max\{|R_w \cap U_{\mathcal{F}}| : w \in W \setminus X\}$ 
10:   $U_{\mathcal{F}} \leftarrow U_{\mathcal{F}} \setminus R_{\bar{w}}$ 
11:   $X \leftarrow X \cup \{\bar{w}\}$ 
12: end while
13: return  $X$ 

```

Note that in the uniform card-V-RAP case, one can start with an arbitrary U -perfect matching M since the cost of any such matching is $|U|$. Furthermore, all job nodes are matched to a vulnerable machine node, as we assume the uniform case. Therefore, each job node is critical and $U_{\mathcal{F}} = U$.

To prove the quality of the computed solution X , we proceed as follows. We distinguish two types of iterations of the set cover greedy subroutine corresponding to steps 8–12 of Algorithm 2. An iteration is called *productive* if the cardinality of $U_{\mathcal{F}}$, the set of not yet covered job nodes, decreases by at least two in step 10 of this iteration. This means that adding the current machine node \bar{w} computed in step 9 to X will saturate at least two critical job nodes. All other iterations are called *nonproductive*.

Let p be the total decrease of $U_{\mathcal{F}}$ obtained from all productive iterations, and let $\text{OPT} \subseteq W$ denote an optimal solution for the given uniform card-V-RAP instance. We next prove two claims that we will use to derive the desired approximation ratio.

• **Claim 1:** $|X| \leq 2|U| - \frac{p}{2}$

In step 2, X is set to W^M , i.e. $|X| = |W^M| = |U|$. In the uniform case, every job node is critical. Thus, the ground set of the set cover instance is U . Since every productive iteration saturates at least two nodes, there can be at most $\frac{p}{2}$ productive iterations. In each such iteration, one further machine node is added to X . All remaining iterations are nonproductive, and there are exactly

$|U| - p$ of them. In total, we have

$$|X| \leq |U| + \frac{p}{2} + (|U| - p) = 2|U| - \frac{p}{2}.$$

- **Claim 2:** $|\text{OPT}| \geq \max\{|U|, 2(|U| - p)\}$

OPT contains at least one U -perfect matching, i.e. $|\text{OPT}| \geq |U|$.

Recall that there are $|U| - p$ nonproductive iterations, and that in each non-productive iteration only one additional job node is covered. We denote by U' all job nodes from U being covered in a nonproductive iteration of the set cover greedy subroutine. Then, for any pair of distinct nodes $u_1, u_2 \in U'$, $u_1 \neq u_2$, we observe that their neighborhoods in G are disjoint i.e. $N_G(u_1) \cap N_G(u_2) = \emptyset$. Indeed, if this were not the case, and u_1 and u_2 would have a common neighbor in W , it would be possible to cover them in a productive iteration, contradicting $u_1, u_2 \in U'$. Since we have a uniform instance, any node in U (and hence in U') must have at least two neighbors in any feasible solution, including OPT . This gives us the bound $|\text{OPT}| \geq 2(|U| - p)$.

Finally, we derive an upper bound on $\frac{|X|}{|\text{OPT}|}$ corresponding to the approximation ratio. Claim 2 allows us to focus on the following two cases.

- **Case 1:** $|\text{OPT}| \geq |U| \geq 2(|U| - p)$, i.e. $2p \geq |U|$

From Claim 1 and $2p \geq |U|$, we obtain that

$$|X| \leq 2|U| - \frac{p}{2} \leq 2|U| - \frac{|U|}{4} \leq \frac{7}{4}|U| \leq \frac{7}{4}|\text{OPT}|.$$

- **Case 2:** $|\text{OPT}| \geq 2(|U| - p) \geq |U|$, i.e. $2p \leq |U|$

From Claim 1 and $2p \leq |U|$, we derive that

$$|X| \leq 2|U| - \frac{p}{2} \leq \frac{7}{2}|U| - \frac{7}{2}p \leq \frac{7}{2}(|U| - p) \leq \frac{7}{4}|\text{OPT}|.$$

In both cases we obtain $\frac{|X|}{|\text{OPT}|} \leq \frac{7}{4}$, which concludes the proof. ■

For non-uniform instances of card-V-RAP the latter proof does not apply, since $|\text{OPT}| \geq 2(|U| - p)$ need not hold in general. However, we can still show that a factor 2 is achieved.

Corollary 8. *Algorithm 2 has an approximation ratio of 2 for non-uniform instances of card-V-RAP.*

Proof. Replace the bounds in Claim 2 with the bounds $|\text{OPT}| \geq |U|$ and $|X| \leq 2|U| - \frac{p}{2}$. Combining these two bounds yields the ratio of 2. ■

4.3 Two vulnerable machines case

One interesting fact about the Edge-Robust Assignment Problem is that this problem is even NP-hard in its simplest non-trivial variant, i.e. in case of two vulnerable edges and unit weights (see [2], and [3, Section 5] for a proof). In contrast, here we show that the node-robust assignment problem remains tractable in this setting. This result is summarized in the next theorem.

Theorem 9. *card-V-RAP with $|\mathcal{F}| = 2$ is solvable in polynomial time.*

Proof. Let $\mathcal{I} = (G, \mathcal{F})$ be a card-V-RAP instance with $G := (U \dot{\cup} W, E)$ and $\mathcal{F} = \{w', w''\} \subseteq W$, $w' \neq w''$. Given an optimal solution X to \mathcal{I} , observe first that either both w' and w'' are contained in X or none of them. In the latter case, an optimal solution is given by an U -perfect matching in $G - \{w', w''\}$. One can use standard bipartite matching algorithms to verify the existence of such a matching and, if existent, corresponds to an optimal solution.

In the remaining case \mathcal{F} is part of any optimal solution. We introduce a dummy job node d and the edges $e' := \{d, w'\}$ and $e'' := \{d, w''\}$. Additionally we double every edge that is not incident to one of the vulnerable nodes w' and w'' . This gives us a new graph $G' := (U' \dot{\cup} W', E')$ with $U' := U \cup \{d\}$, $W' = W$ and $E' := E \cup \{e', e''\} \cup \{\bar{e} \mid e \in E \setminus \delta_G(\{w', w''\})\}$ (see Figure 4, for an illustration). Note that




Figure 4: Input graph G (left), graph G' resulting by introducing the node d and the edges e' and e'' (center) and a solution to the corresponding ILP (6) (right).

the new graph G' remains bipartite.

We will obtain an optimal solution for the V-RAP instance \mathcal{I} from the following ILP.

$$\begin{aligned}
\min \quad & \sum_{e \in E'} x_e \\
\text{s.t.} \quad & x(\delta(u)) = 2, \text{ for each } u \in U', \\
& x_{e'} = x_{e''} = 1, \\
& x \in \{0, 1\}^{E'}.
\end{aligned} \tag{6}$$

Every solution to ILP (6) forms a collection of cycles of size greater than or equal to two that cover every node from U' . Cycles of size two are formed by parallel edges in G' , i.e. each such cycle represents an original edge from G . The cycle covering node d contains the newly introduced edges e' and e'' , and has a size of at least four. This cycle corresponds to a path from w' to w'' in G with an even number of edges. Thus, every solution to ILP (6) defines a union of a w' - w'' -path, a (possibly empty) matching, some additional even paths, and potentially some further cycles in the original graph G . As G' is bipartite, the constraint matrix of ILP (6) is totally unimodular (see [16] for background on totally unimodular matrices) implying that ILP (6) can be solved in polynomial time via LP methods.

Next, we claim that solutions to ILP (6) correspond to inclusion-wise minimal solutions to the V-RAP instance \mathcal{I} and vice versa. For a given feasible solution x of ILP (6), consider

$$X := \{w \in W \mid \exists e \in \delta(w) \text{ with } x_e = 1\}.$$

We argue that X is feasible to the original card-V-RAP instance \mathcal{I} . To see this, note that each original job node u is adjacent to at least one non-vulnerable machine node from X . Each node $u \in U$ located on a cycle in G' that is induced by x and does not contain a vulnerable node can be easily matched in $G[U \dot{\cup} X]$ as u is either incident to an isolated edge in $G[U \dot{\cup} X]$ or is part of an even cycle. Even paths in G' correspond to even paths in G . Since job nodes are inner nodes of these paths, they can be matched using edges from these paths. All remaining nodes from U are inner nodes of the even w' - w'' -path induced by x , i.e. we can find a U -perfect matching in $G[U \dot{\cup} X]$ not using w' and w'' , simultaneously. This shows that any feasible solution x of ILP (6) corresponds to a solution X feasible to \mathcal{I} .

Now, let X be any feasible solution to \mathcal{I} . Then, $G[U \dot{\cup} X]$ contains a matching M' with $w' \notin V(M')$ and a matching M'' with $w'' \notin V(M'')$. By our assumption that $\{w', w''\}$ must be contained in every feasible solution for \mathcal{I} , we have that $w'' \in V(M')$ and $w' \in V(M'')$. Then, the symmetric difference $M' \Delta M''$ contains an even path P from w' to w'' . Moreover, $\hat{M} := M' \setminus E(P)$ is a possibly empty matching on the job nodes not covered by the path P . Consider $x \in \mathbf{R}^{E'}$ with

$$x_e = \begin{cases} 1, & \text{if } e \in \{e', e''\} \cup E(P), \\ 1, & \text{if } e \in \hat{M} \cup \{\bar{e} \mid e \in \hat{M}\}, \\ 0, & \text{otherwise.} \end{cases}$$

By construction the vector x satisfies the constraints of ILP (6).

Finally, solutions of ILP (6) can be used to obtain optimal solutions to \mathcal{I} as follows. In the graph $G[U \dot{\cup} X]$ a job node can be adjacent to two non-vulnerable machines (see Figure 4). Such solutions are not optimal and can be identified by checking if $G[U \dot{\cup} X]$ has any odd component not containing w' and w'' . In each of these components the number of machine nodes exceeds the number of job nodes by one. Now by removing an arbitrary machine node from each component we obtain an optimal solution to \mathcal{I} . ■

Acknowledgement. The work of the second and the third author is part of the Research Training Group “Discrete Optimization of Technical Systems under Uncertainty”. Financial support by the German Research Foundation (DFG) is gratefully acknowledged through RTG 1855. Parts of this work were carried out during the second author was visiting IFOR at ETH Zurich. The hospitality is highly appreciated.

5 References

- [1] David Adjiashvili, Sebastian Stiller, and Rico Zenklusen. Bulk-robust combinatorial optimization. *Mathematical Programming*, 149(1-2):361–390, 2015. URL <http://dx.doi.org/10.1007/s10107-014-0760-6>. 2
- [2] David Adjiashvili, Viktor Bindewald, and Dennis Michaels. Robust Assignments via Ear Decompositions and Randomized Rounding. In *ICALP 2016*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 71:1–71:14, Dagstuhl, Germany, 2016. ISBN 978-3-95977-013-2. URL <http://drops.dagstuhl.de/opus/volltexte/2016/6213>. 1, 2, 5, 6, 8, 19, 21
- [3] David Adjiashvili, Viktor Bindewald, and Dennis Michaels. Robust assignments via ear decompositions and randomized rounding. *ArXiv e-prints*, 2016. URL <http://arxiv.org/abs/1607.02437>. full preprint version of [2]. 2, 6, 7, 8, 19
- [4] Robert E. L. Aldred, Richard P. Anstee, and Stephen C. Locke. Perfect matchings after vertex deletions. *Discrete Mathematics*, 307(23):3048–3054, 2007. ISSN 0012-365X. URL <http://www.sciencedirect.com/science/article/pii/S0012365X07001069>. 8
- [5] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In *Algorithms and Complexity*, pages 288–298. Springer, 1997. URL http://link.springer.com/chapter/10.1007/3-540-62592-5_80. 14, 16

- [6] Ashwin Arulselvan, Ágnes Cseh, Martin Groß, David F. Manlove, and Jan-nik Matuschke. Many-to-one matchings with lower quotas: Algorithms and complexity. In *ISAAC 2015 Proceedings*, pages 176–187, Berlin, Heidelberg, 2015. ISBN 978-3-662-48971-0. URL http://dx.doi.org/10.1007/978-3-662-48971-0_16. 7
- [7] Vašek Chvátal. A Greedy Heuristic for the Set-Covering Problem. *Mathematics of Operations Research*, 4(3):233–235, 1979. ISSN 0364-765X. URL <http://pubsonline.informs.org/doi/abs/10.1287/moor.4.3.233>. 8, 14
- [8] Mitre C. Dourado, Dirk Meierling, Lucia D. Penso, Dieter Rautenbach, Fabio Protti, and Aline Ribeiro de Almeida. Robust recoverable perfect matchings. *Networks*, 66(3):210–213, 2015. ISSN 1097-0037. URL <http://dx.doi.org/10.1002/net.21624>. 7
- [9] Andreas T. Ernst, Houyuan Jiang, Mohan Krishnamoorthy, and David Sier. Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research*, 153(1):3–27, 2004. ISSN 0377-2217. URL <http://www.sciencedirect.com/science/article/pii/S037722170300095X>. 3
- [10] Odile Favaron. On k -factor-critical graphs. *Discuss. Math. Graph Theory*, 16(1):41–51, 1996. ISSN 1234-3099. URL <http://dx.doi.org/10.7151/dmgt.1022>. 8
- [11] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998. URL <http://doi.acm.org/10.1145/285055.285059>. 11
- [12] Irit Katriel, Claire Kenyon-Mathieu, and Eli Upfal. Commitment under uncertainty: Two-stage stochastic matching problems. *Theoretical Computer Science*, 408(2):213–223, 2008. ISSN 0304-3975. URL <http://www.sciencedirect.com/science/article/pii/S0304397508005641>. 7
- [13] Mathieu Lacroix, Ali R. Mahjoub, Sébastien Martin, and Christophe Picouleau. On the np-completeness of the perfect matching free subgraph problem. *Theoretical Computer Science*, 423:25 – 29, 2012. ISSN 0304-3975. URL <http://www.sciencedirect.com/science/article/pii/S0304397511010413>. 7
- [14] Pierre Laroche, Franc Marchetti, Sébastien Martin, and Zsuzsanna Róka. Bipartite complete matching vertex interdiction problem: Application to robust nurse assignment. In *Control, Decision and Information Technologies (CoDIT), 2014 International Conference on*, pages 182–187, 2014. URL <http://dx.doi.org/10.1109/CoDIT.2014.6996890>. 7

- [15] László Lovász and Michael D Plummer. *Matching theory, volume 121 of North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1986. 13
- [16] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics. Springer Berlin Heidelberg, 2002. ISBN 9783540443896. URL <http://books.google.de/books?id=19rCngEACAAJ>. 20
- [17] Rico Zenklusen. Matching interdiction. *Discrete Applied Mathematics*, 158(15): 1676 – 1690, 2010. URL <http://www.sciencedirect.com/science/article/pii/S0166218X10002143>. 7